



Applied Micro Circuits Corporation
***Challenges in Making Highly
Integrated Network Processors***

Keith Morris
Director of Marketing
Switching & Network Processing
kmorris@amcc.com

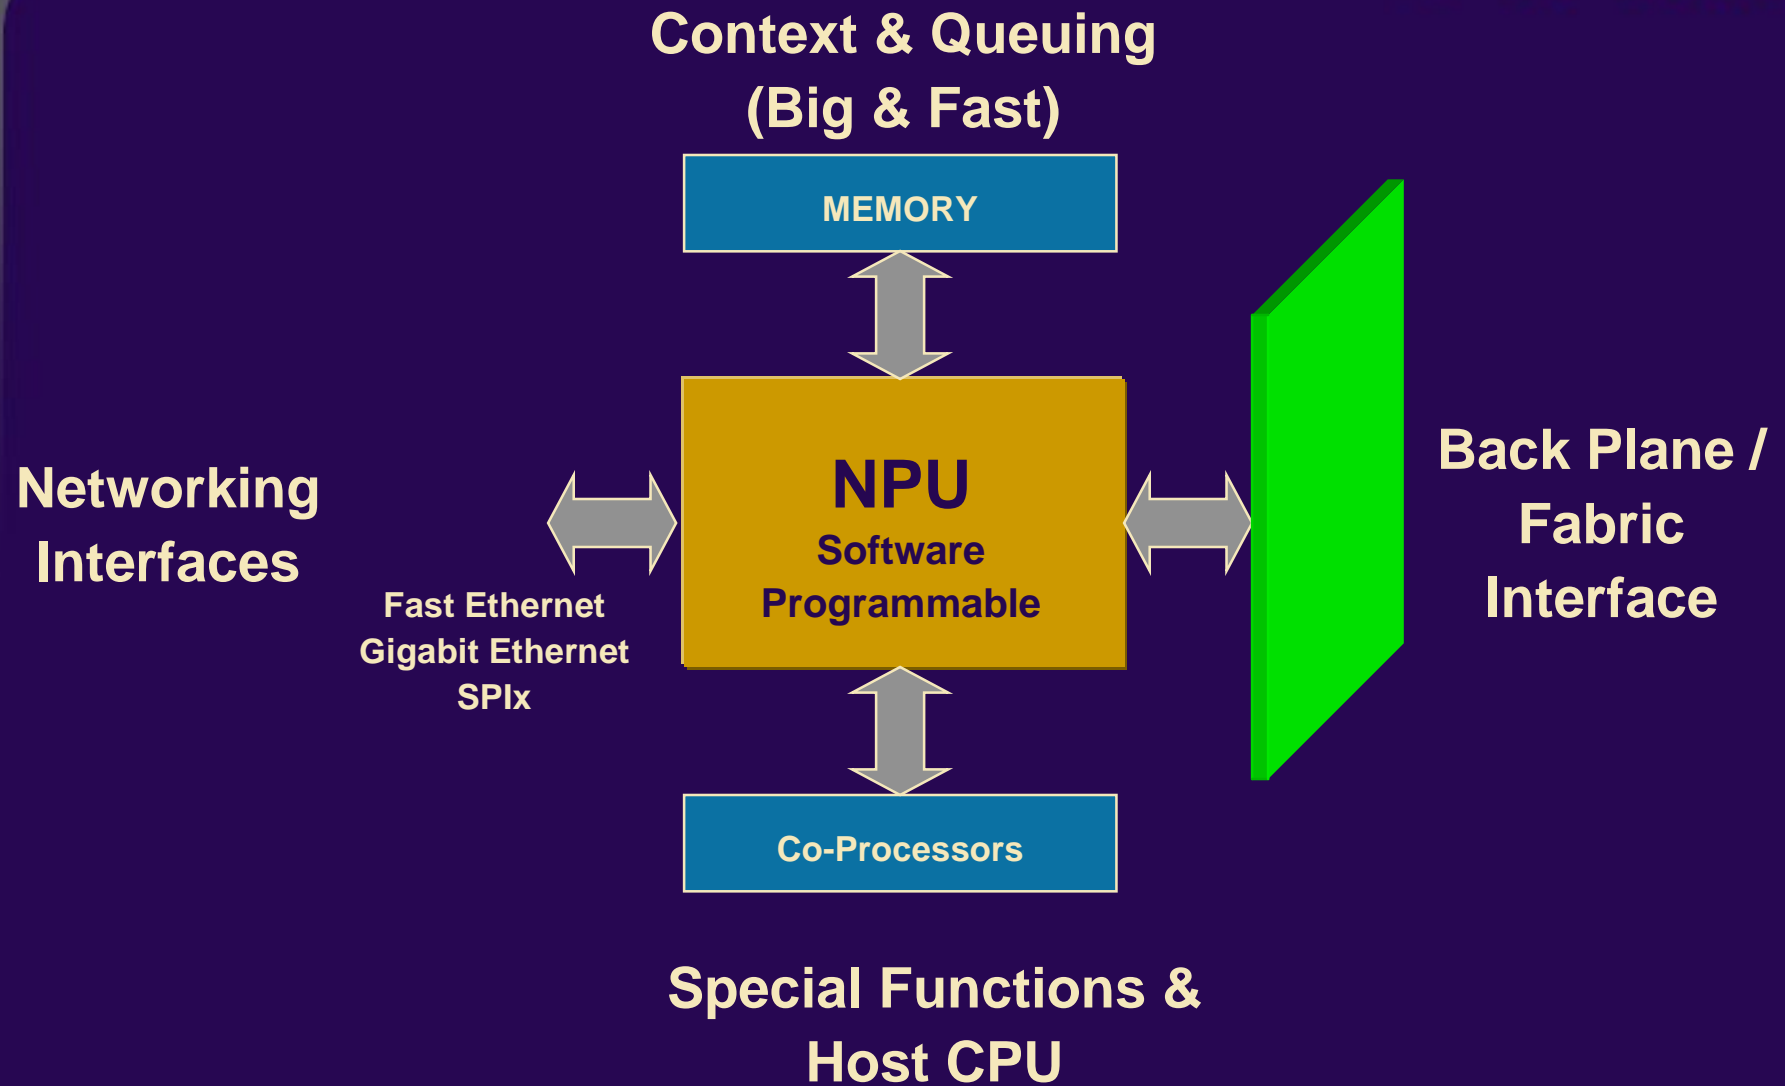
CONFIDENTIAL



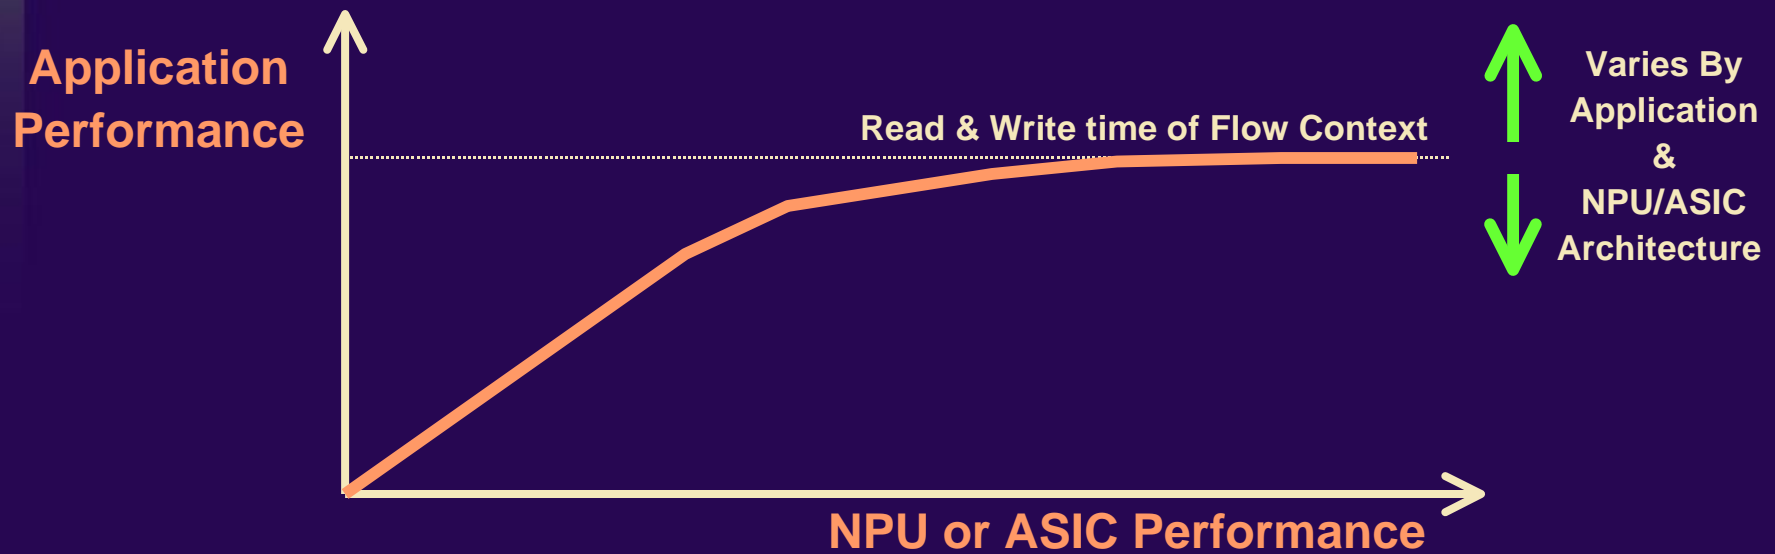
Agenda

- **NPU Definition**
- **Challenges & Solutions**
 - **Physical**
 - **Performance**
 - **Flexibility**
 - **Scalability & Re-use**
- **Putting it all together – Integrated NPU**

NPU Definition



Context Handling – The Ultimate Bottleneck






Challenges & Solutions

CONFIDENTIAL





Challenge: Physical Constraints




Small Silicon Die Area

Packet Arrival Rate



System Power Dissipation



Efficient nPcore Architecture

Networking Optimized Execution Pipeline


Overlapping Virtual Processors

T0	T1	T2	T3	T4	T5
				T0	T1
Task Sel	Task Sel	Task Sel	Task Sel	Task Sel	Task Sel
	Fetch	Fetch	Fetch	Fetch	Fetch
		Decode	Decode	Decode	Decode
			Read	Read	Read
				Execute	Execute
					Wr Back


Case Statements and Conditional Jumps do not cause pipeline breakages

Extremely important for deterministic performance in decision rich data path processing

Zero Cycle Task Switching – Hides Latency




Frequency, Power and Die Area for nPcores



Challenge: Scaling Processing Performance


Processor



Pipeline – Multi-Stage



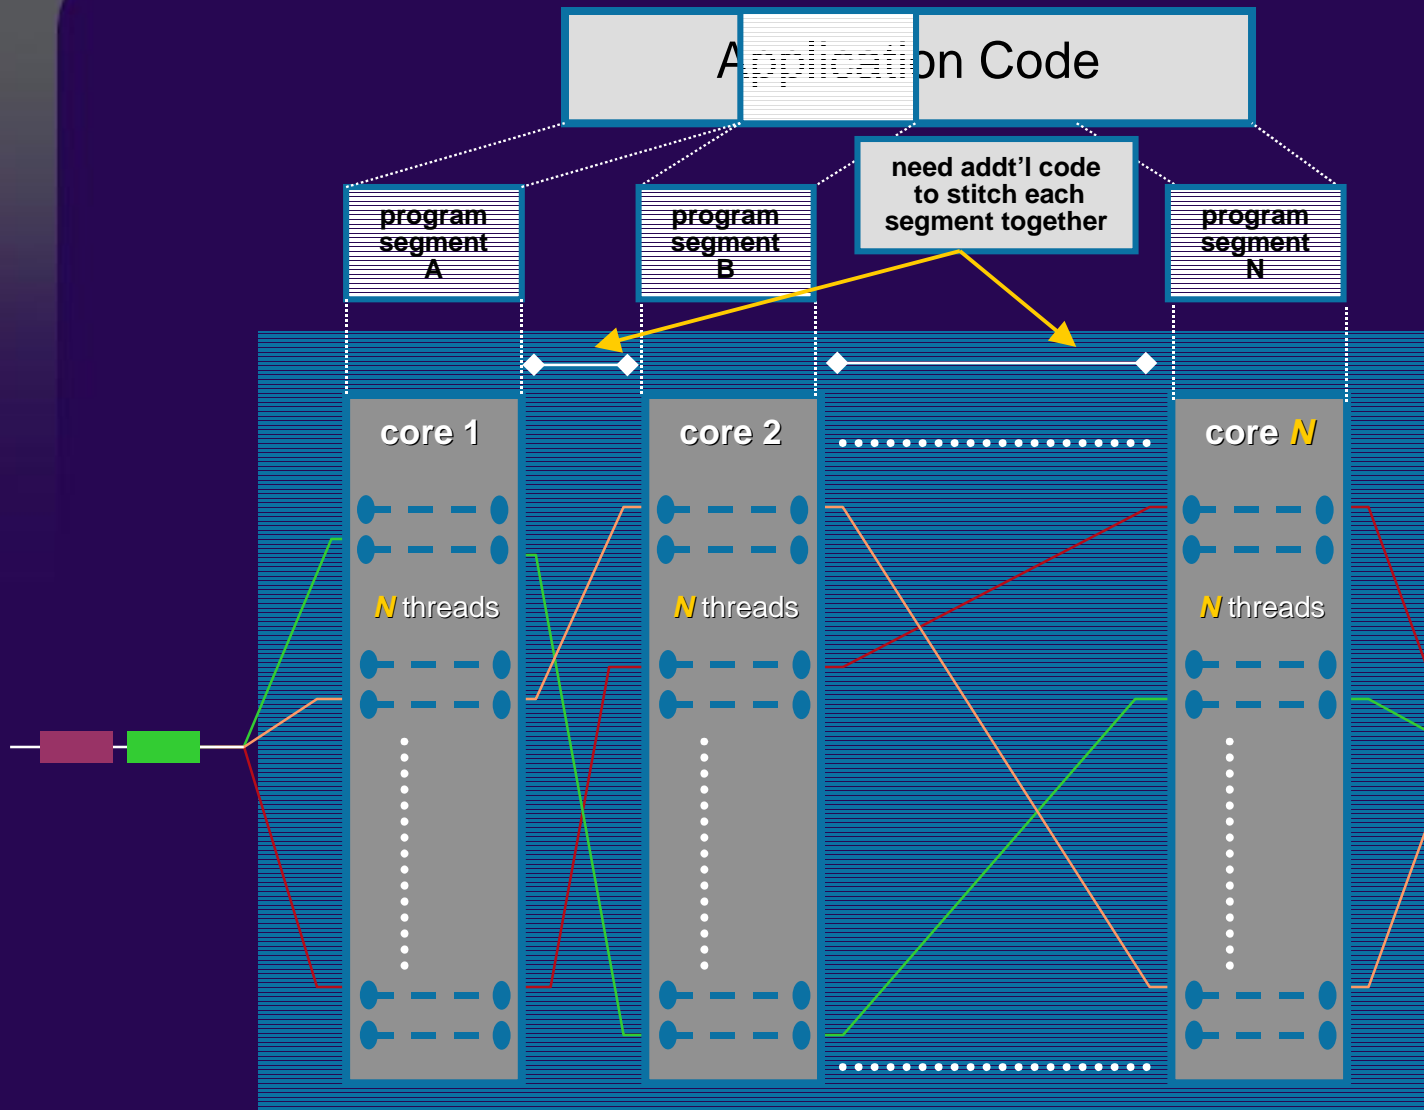
- Complex software partitioning
- Very sensitive to latency
 - Each stage must complete within the smallest packet time
 - One stage becomes the bottleneck
- Poor Scalability
 - Eventually some stages may need to become parallel



Parallel – Single Stage

- Allows simplest programming model
 - Run to completion
- Not sensitive to latency
 - Elapsed time to process a packet can be very long if required
- Highly scalable
 - 100's of tasks, multiple processors

Multi-Stage – Complex and Inefficient



Developer must:

- Subdivide algorithm
- Load-balance segments
- Stitch segments together

Performance issues with:

- Underutilized cores
- 1 segment can overrun next
- Hand-offs between cores
- Locking shared resources

AMCC: 1-Stage – Fundamentally Simpler




- Same program image loaded on each thread
- A packet runs to completion on a single thread

```
;; cannot handle.  
;; An encapsulation header is used to communicate the  
;; source-port-id and the errno value as follows:  
;;  
;;  
;;


|    |    |        |       |     |   |   |   |
|----|----|--------|-------|-----|---|---|---|
| 0  | 1  | 2      | 3     | 4   | 5 | 6 | 7 |
| 81 | 00 | S-PORT | ERRNO | N/A |   |   |   |

  
;;  
;;  
=====
```

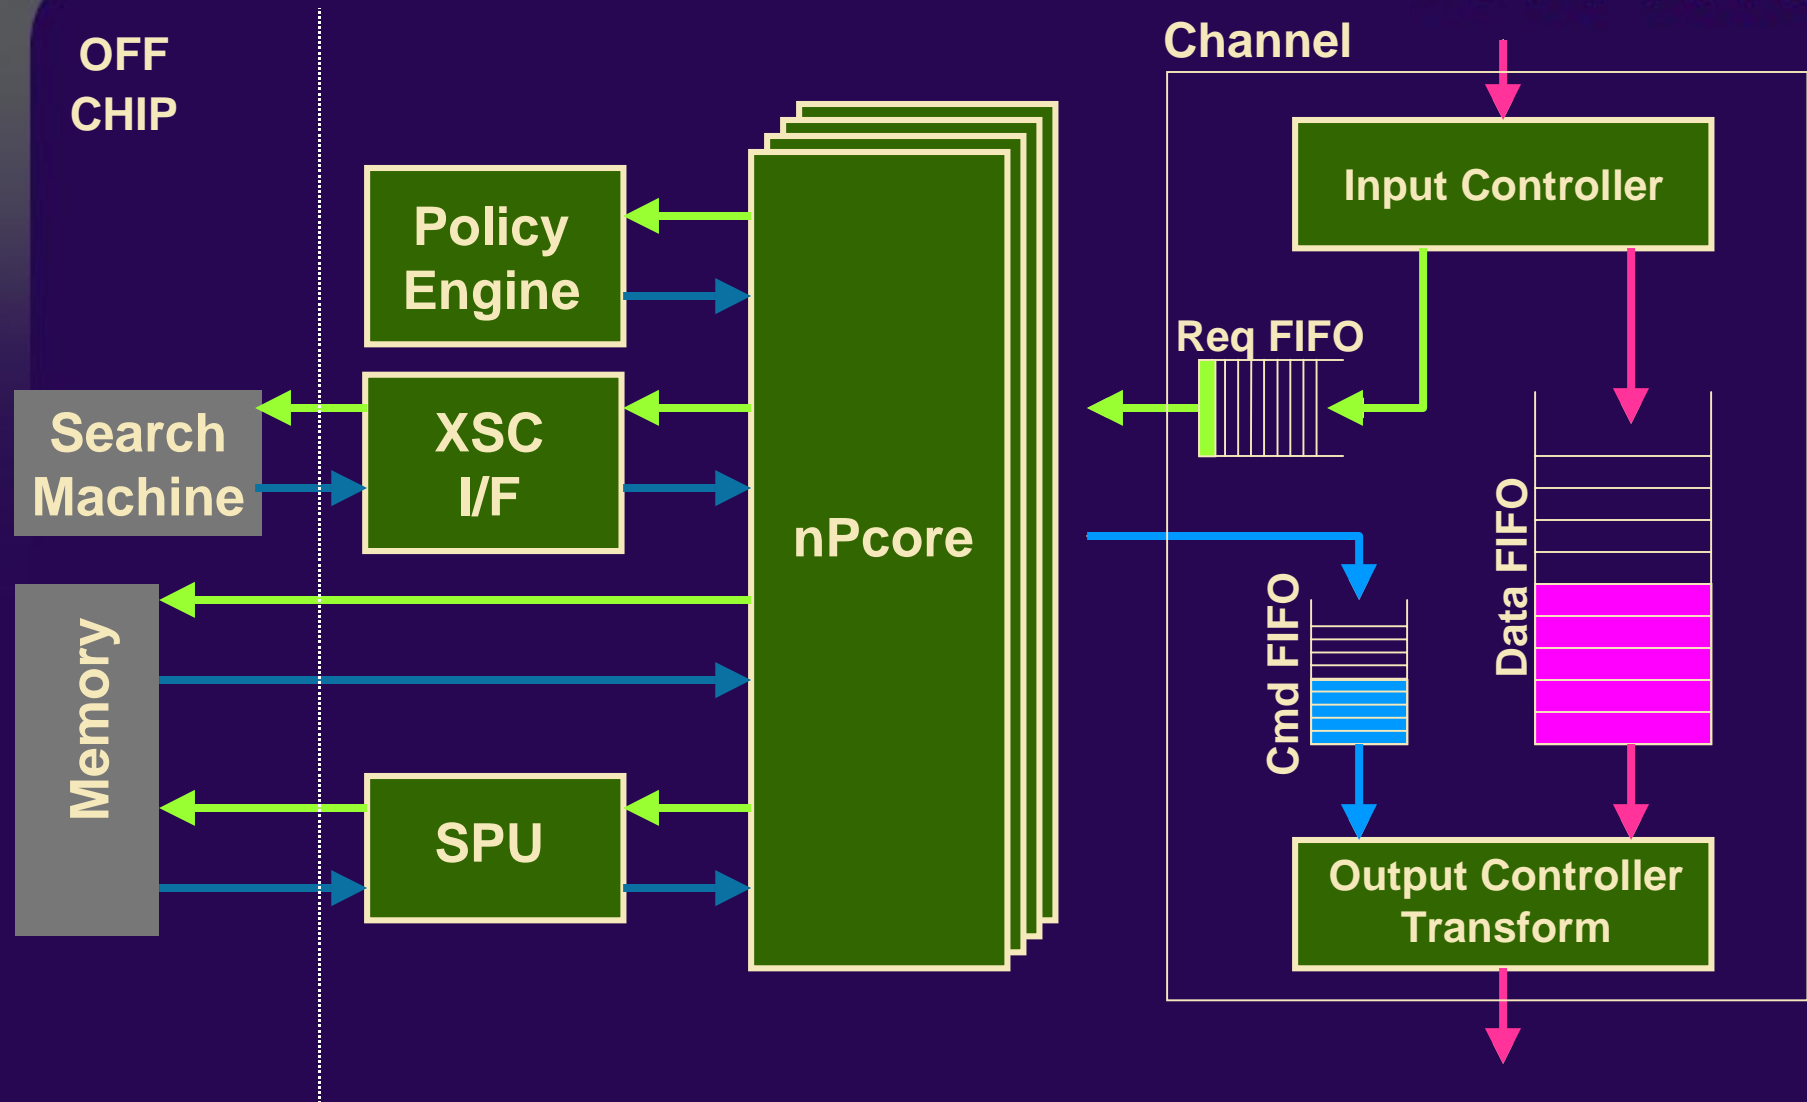
```
pkt_fwd_to_cpu  
#define DB_HDL r6  
nPkernel ;; retrieve the control connection switch header from the  
;; unicast port table.  
set_entry db ucast_port_table, DB_HDL, APP_CPU_PORTID  
read_first_db ucast_port_table, DB_HDL, r0  
;; setup the encap header
```



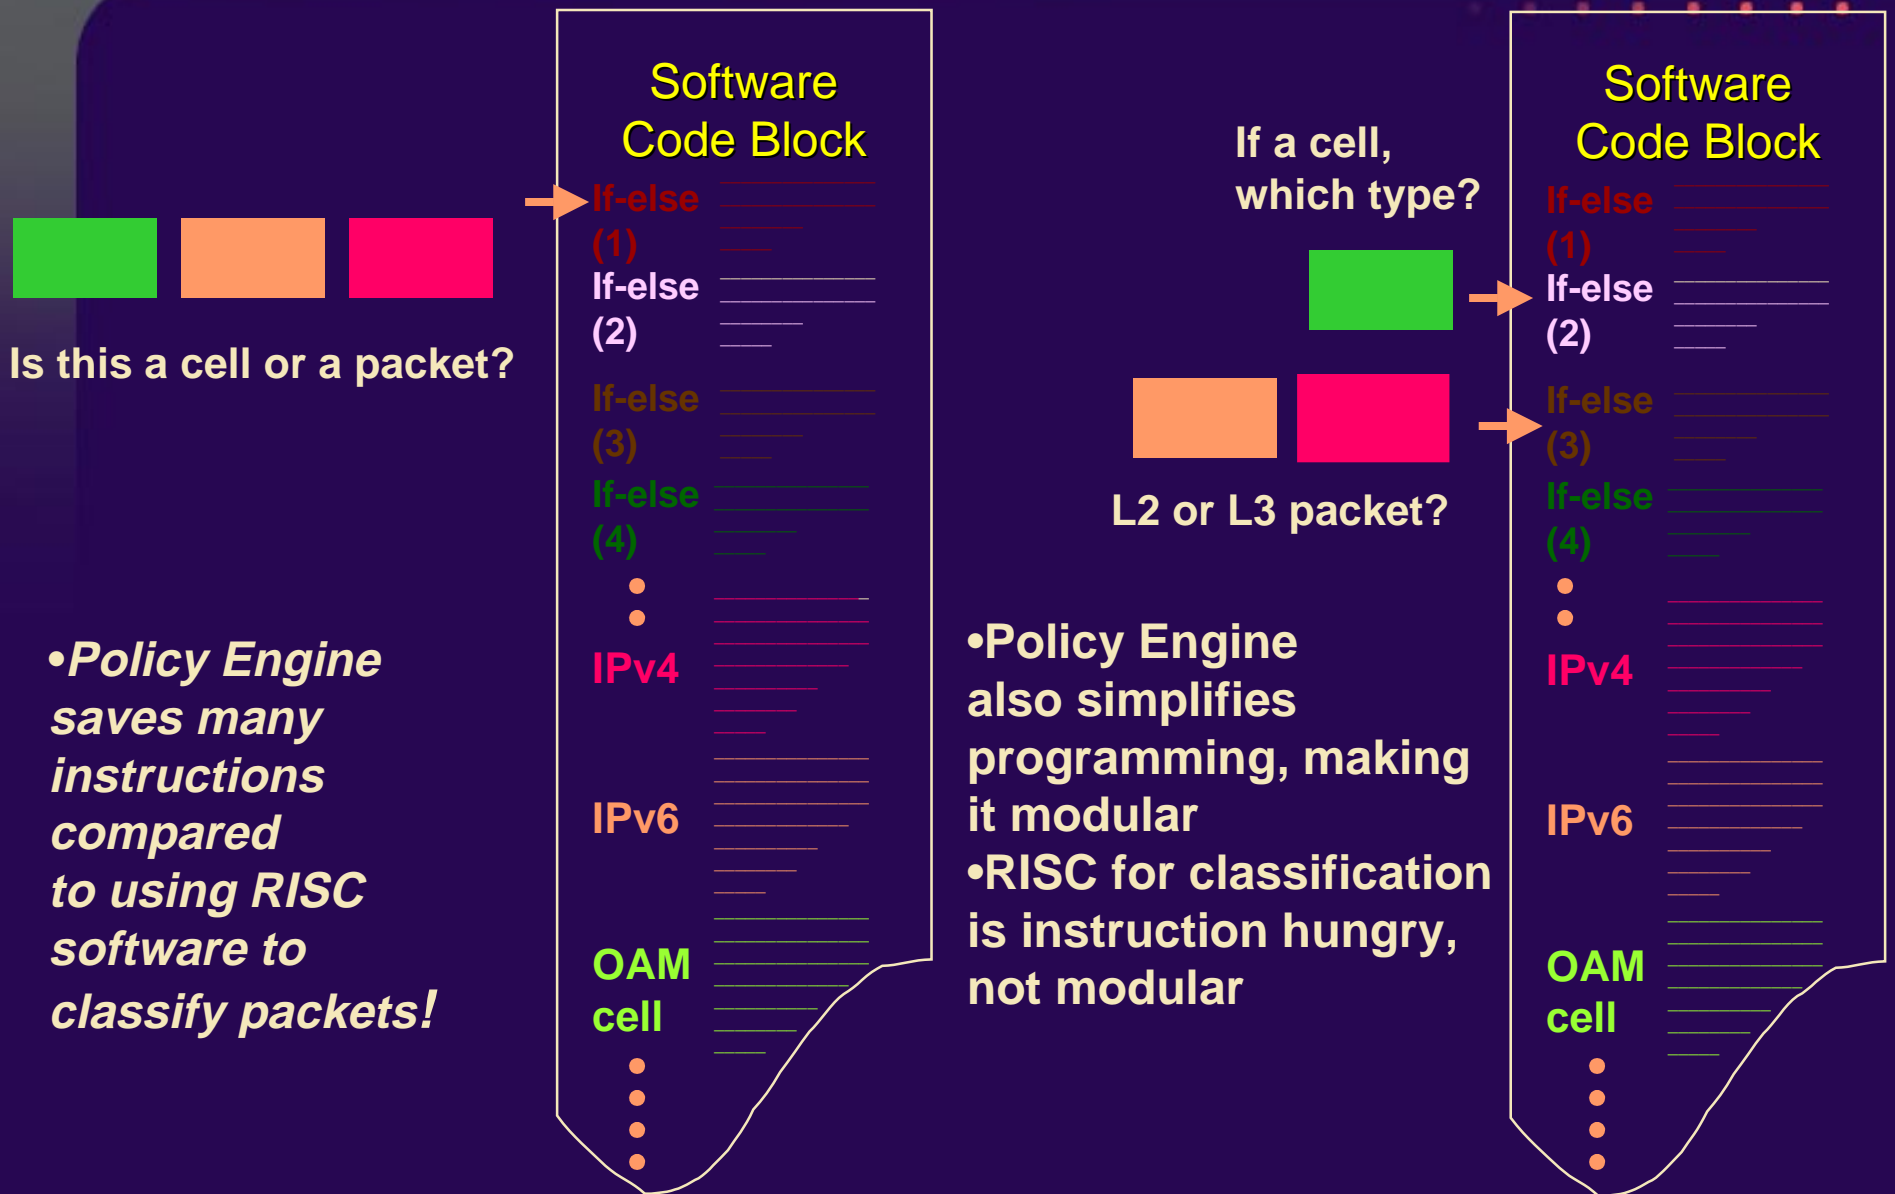
Challenge: Flexibility vs. Performance




Solution: Co-Processors



High Speed Single-Step Packet Classification



Efficient Packet Modification & Encapsulation



Complex Bandwidth Provisioning

PHYSICAL PORT

- OC-192

SUB-PORTS


- Fixed sub-division of the physical port
- OC-192, -48, -12, -3, GE
- Flows are scheduled according to minimum rate, class of service and weight at this level

VIRTUAL PIPES


- A collection of flows that have an aggregate maximum rate
- Pipe ensures that provisioned rate is not exceeded
- For example - all flows for an individual subscriber, network or traffic type

FLOWS

- Smallest scheduled entity
- May have guaranteed minimum bandwidth
- Individually weighted within a pipe

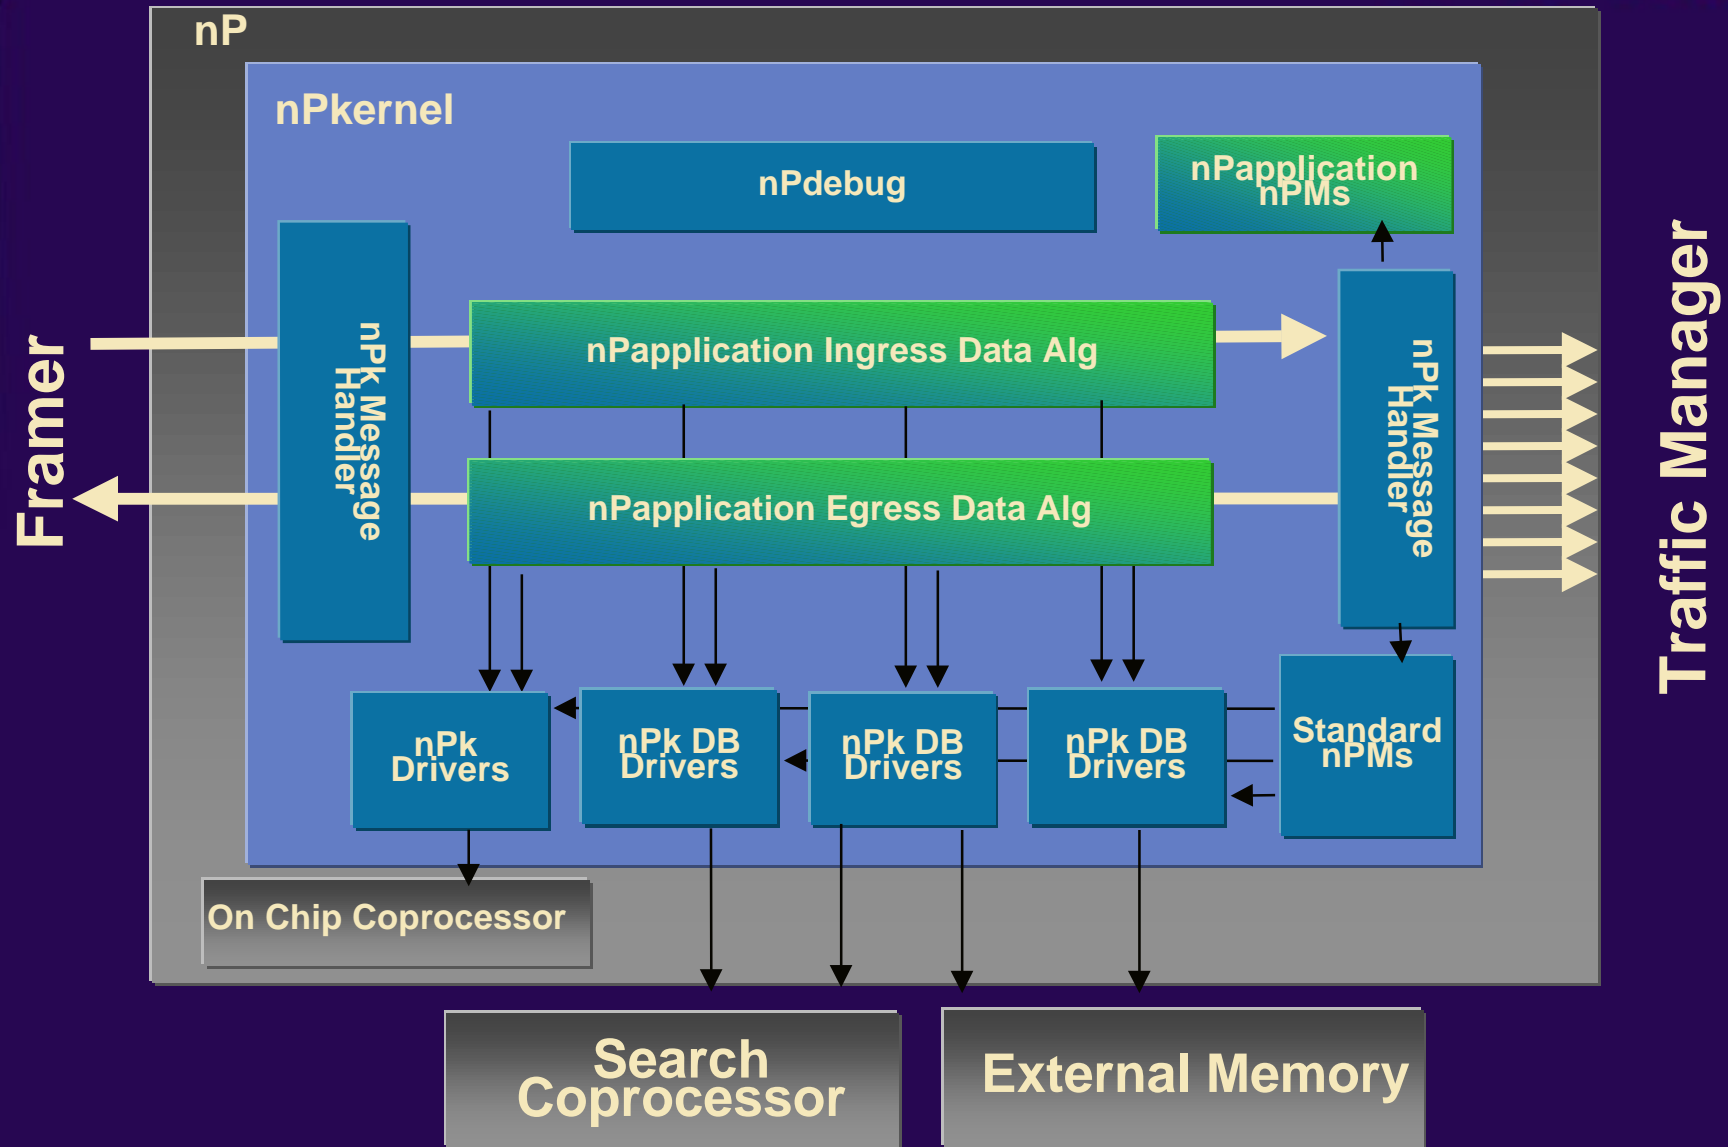


Challenge: Scalability & Re-use




Maximum Software ROI

Solution: nP Programming Infrastructure

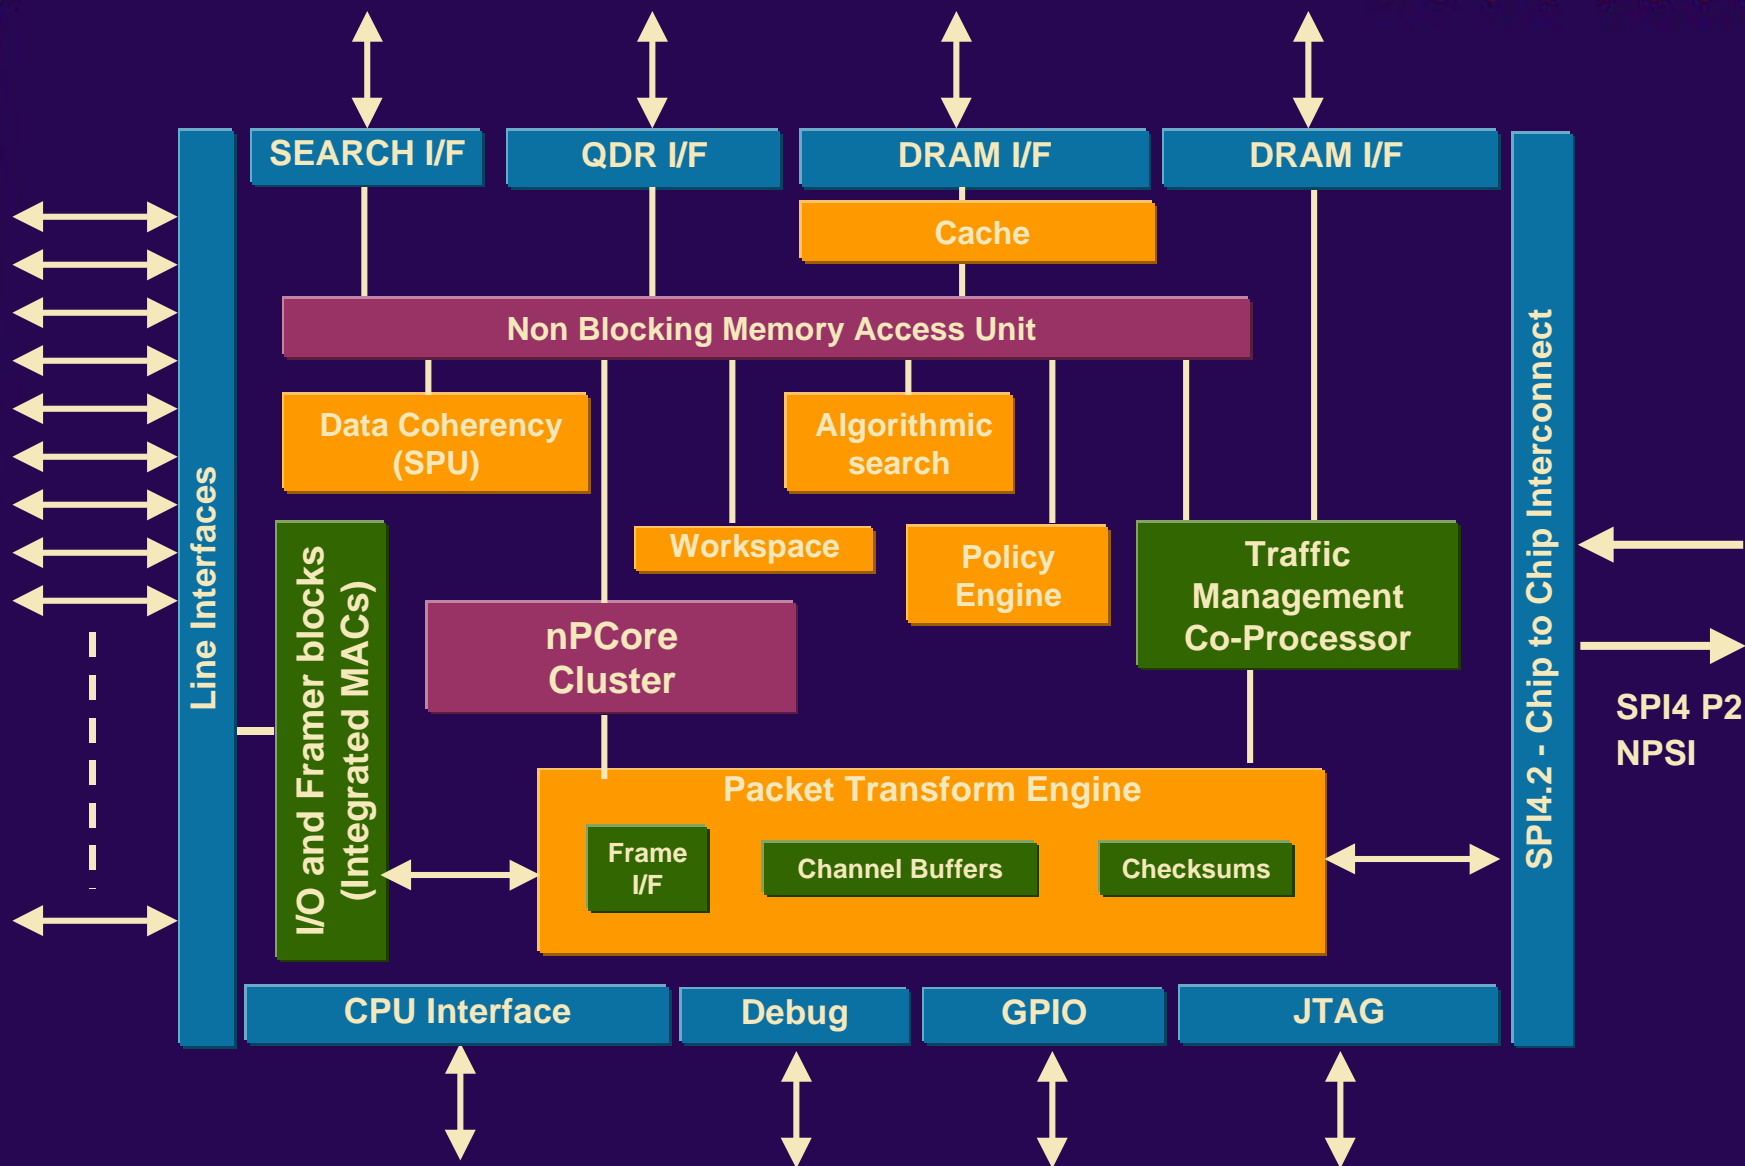


“Runtime Libraries”: Reducing code customer must actually write

- 99+% of total Lines of Code (LOC) are on the control CPU
- Next, nPkernel alone provides up to 95% of actual NPU code, i.e. 95% of the 1%
- So NPU-resident portion of app often only 100-200 total LOC
- AMCC sample app u-code provides most of that 100-200 LOC, all in some cases
- Actual customer-generated code typically much less than 5% of 1% !!



Putting it all together - Integrated NPU



Summary

- **NPU Definition – No performance penalty**
- **Challenges & Solutions**
 - **Physical – More MHz won't cut it**
 - **Performance – Cannot be at the price of usability**
 - **Flexibility – Doesn't mean 100% software**
 - **Scalability & Re-use – Any Protocol, Any Speed, One Architecture**
- **Putting it all together**
 - **Avoiding System Problems & Side Effects**